

Exhibit 3







Exhibit 2

U.S. Patent No. 7,519,814 vs. Oracle

Accused Instrumentalities: Oracle products and services using secure containerized applications, including without limitation Oracle Cloud Infrastructure (“OCI”) and Oracle Kubernetes Engine (“OCE”), and all versions and variations thereof since the issuance of the asserted patent.

Claim 1

Claim 1	Accused Instrumentalities
<p>[1pre] 1. In a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service, the method comprising:</p>	<p>To the extent the preamble is limiting, Oracle and/or its customer practices, through the Accused Instrumentalities, in a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service, as claimed.</p> <p>For example, Oracle Kubernetes Engine runs on individual servers, each of which runs an independent operating system running either on bare metal, through an on-premises virtualized infrastructure, through one or more cloud services, or through any other supported deployment. In an exemplary deployment, two or more servers use different operating systems. The servers operate in disparate computing environments, including because each server is a stand-alone computer and/or each server is unrelated to the other servers due to having independent hardware and, in some instances, independent software.</p> <p>Oracle requires and/or provides that each server includes a processor with one or more cores available to the OS kernel. Oracle further requires and/or provides that each server has a supported operating system, which includes a kernel and associated local system files, including for example libraries such as libc/glibc, configuration files, etc. In the infringing system, at least two servers have different operating systems.</p> <p>In at least some instances, Oracle directly owns, operates, controls, and/or benefits from the claimed system and/or method. In other instances, Oracle’s customer makes and uses the system and/or method either by following Oracle’s direction and control, including Oracle’s documentation, or automatically through the ordinary and expected operation of Oracle’s software, or a combination thereof.</p>

Claim 1	Accused Instrumentalities
	<p>See claim limitations below.</p> <p>See also, e.g.:</p> <h2 data-bbox="699 329 1094 378">Why Choose OKE?</h2> <div data-bbox="699 451 1942 1141"> <div data-bbox="699 451 1081 768">  <p>Price-Performance</p> <p>OKE is the lowest cost Kubernetes service amongst all hyperscalers, especially for serverless.</p> </div> <div data-bbox="1129 451 1530 768">  <p>Autoscaling</p> <p>OKE automatically adjusts compute resources based on demand, which can reduce your costs.</p> </div> <div data-bbox="1570 451 1942 768">  <p>Efficiency</p> <p>GPUs can be scarce, but OKE job scheduling makes it easy to maximize resource utilization.</p> </div> <div data-bbox="699 841 1081 1141">  <p>Portability</p> <p>OKE is consistent across clouds and on-premises, enabling portability and avoiding vendor lock-in.</p> </div> <div data-bbox="1129 841 1530 1141">  <p>Simplicity</p> <p>OKE reduces the time and cost needed to manage the complexities of Kubernetes infrastructure.</p> </div> <div data-bbox="1570 841 1942 1141">  <p>Reliability</p> <p>Automatic upgrades and security patching boost reliability for the control plane and worker nodes.</p> </div> </div> <p data-bbox="674 1187 1478 1219">https://www.oracle.com/cloud/cloud-native/kubernetes-engine/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 215 1696 269">Welcome to Oracle Cloud Infrastructure</h2> <p data-bbox="695 318 1944 467">Oracle Cloud Infrastructure (OCI) is a set of complementary cloud services that enable you to build and run a range of applications and services in a highly available hosted environment. OCI provides high-performance compute capabilities (as physical hardware instances) and storage capacity in a flexible overlay virtual network that is securely accessible from your on-premises network.</p> <p data-bbox="674 492 1667 521">https://docs.oracle.com/en-us/iaas/Content/GSG/Concepts/baremetalintro.htm</p> <h3 data-bbox="722 581 1871 695">Existing applications can benefit by migrating to OCI and OKE</h3> <p data-bbox="722 743 1797 781">OKE offers lower total cost of ownership and improved time to market.</p> <p data-bbox="722 846 1572 883">OKE simplifies operations at scale in the following ways:</p> <ul data-bbox="722 967 1871 1344" style="list-style-type: none"> <li data-bbox="722 967 1247 1052">– Lift and shift; there’s no need to rearchitect <li data-bbox="1325 967 1850 1052">– Increase resource utilization and efficiency <li data-bbox="722 1117 1247 1201">– Reduce operations burden with automation <li data-bbox="1325 1117 1850 1201">– Improve agility, flexibility, uptime, and resilience <li data-bbox="722 1266 1247 1351">– Save time on infrastructure management <li data-bbox="1325 1266 1850 1351">– Reduce compliance risk and enhance security <p data-bbox="674 1393 1675 1422">https://www.oracle.com/cloud/cloud-native/kubernetes-engine/#app-migration</p>

Claim 1	Accused Instrumentalities
	<p>What is OCI Kubernetes Engine (OKE)?</p> <p>Oracle Cloud Infrastructure Kubernetes Engine (OKE) is a managed Kubernetes service that simplifies the development, deployment, and operation of containerized workloads at scale. OKE enables you to quickly create, manage, and consume Kubernetes clusters that leverage underlying OCI compute, networking, and storage services.</p> <p>When should I use OKE?</p> <p>You should use OKE when you want to leverage Kubernetes to deploy and manage your Kubernetes-based container applications. It allows you to combine the production-grade container orchestration of standard upstream Kubernetes with the control, security, and high, predictable performance of OCI.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <p>How does OKE provide resiliency?</p> <p>When you create an OKE cluster, OKE automatically creates and manages multiple Kubernetes control plane nodes spread across fault domains and availability domains (logical data centers). This is done to help ensure that the managed Kubernetes control plane is highly available. Control plane operations, such as upgrading to newer versions of Kubernetes, can be performed without service interruptions. Additionally, when you provision worker nodes, you can use a placement configuration to control the fault domain and availability domain where they are created. Nodes will automatically come online with labels, which you can use to schedule your workloads so they are robust and highly available.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <p>Can I deploy private Kubernetes clusters?</p> <p>Yes; with OKE, your Kubernetes clusters are integrated in your VCN. Your cluster worker nodes, load balancers, and the Kubernetes API endpoint are part of a private or public subnet of your VCN. Regular VCN routing and firewall rules control access to the Kubernetes API endpoint, making it accessible from a corporate network only, through a bastion host, or by specific platform services.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p>

Claim 1	Accused Instrumentalities
	<p>When should I use virtual nodes, managed nodes, or self-managed nodes?</p> <ul style="list-style-type: none"> Virtual nodes Virtual nodes offer a serverless Kubernetes experience. This option is ideal if you'd rather focus on your application and avoid managing the underlying infrastructure. Virtual nodes relieve you of management-related tasks such as scaling, upgrading, patching, troubleshooting, and provisioning worker nodes. Managed nodes Managed nodes are a good choice for general purpose workloads. They offer an extensive list of customizable configuration options that have been tested by the OKE service. Unlike fully managed virtual nodes, you share the management of worker nodes with OCI. OKE simplifies the management process through features such as on-demand cycling to automate worker node updates, cluster self-healing upon failure detection, autoscaling, and more. Self-managed nodes Self-managed nodes offer access to the underlying infrastructure, configuration options, and compute shapes that aren't currently available to managed nodes. This includes access to specialized infrastructure, such as RDMA-enabled bare metal cluster networks or confidential compute shapes. This advanced control makes self-managed nodes ideal for specialized use cases that aren't supported with managed nodes. Note that with self-managed nodes, you are fully responsible for managing the worker nodes—without the automated features provided by managed or virtual nodes. <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p> <p>What are the storage options for virtual nodes?</p> <p>OKE virtual nodes do not yet have persistent storage capabilities. However, there are plans to introduce support for attaching persistent volumes backed by OCI Block Storage and OCI File Storage. If your Kubernetes application requires persistent storage, it's advisable to use OKE managed nodes.</p> <p>https://www.oracle.com/cloud/cloud-native/kubernetes-engine/faq/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="711 215 1581 264">Supported Images for Managed Nodes</h2> <p data-bbox="711 313 1959 378">Kubernetes Engine supports the provisioning of worker nodes (managed nodes only) using some, but not all, of the latest Oracle Linux images provided by Oracle Cloud Infrastructure.</p> <p data-bbox="711 427 1602 451">You can use these Oracle Linux images when provisioning managed nodes:</p> <ul data-bbox="737 500 972 638" style="list-style-type: none"> • OKE Images • Platform Images • Custom Images <p data-bbox="674 670 1818 703">https://docs.oracle.com/en-us/iaas/Content/ContEng/Reference/contengimagesshapes.htm</p> <h2 data-bbox="684 743 1014 784">What is Docker?</h2> <hr data-bbox="684 824 747 833" style="width: 30px; margin-left: 0;"/> <p data-bbox="684 898 1944 1052">A Docker container is a packaging format that packages all the code and dependencies of an application in a standard format that allows it to run quickly and reliably across computing environments. A Docker container is a popular lightweight, standalone, executable container that includes everything needed to run an application, including libraries, system tools, code, and runtime. Docker is also a software platform that allows developers to build, test, and deploy containerized applications quickly.</p> <p data-bbox="684 1084 1959 1182">Containers as a Service (CaaS) or Container Services are managed cloud services that manage the lifecycle of containers. Container services help orchestrate (start, stop, scale) the runtime of containers. Using container services, you can simplify, automate, and accelerate your application development and deployment lifecycle.</p> <p data-bbox="684 1214 1902 1304">Docker and Container Services have seen rapid adoption and have been a tremendous success over the last several years. From an almost unknown and rather technical open source technology in 2013, Docker has evolved into a standardized runtime environment now officially supported for many Oracle enterprise products.</p> <p data-bbox="674 1336 1707 1360">https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>Docker versus Kubernetes</p> <p>Linux containers have existed since 2008, but they were not well known until the emergence of Docker containers in 2013. With the onset of Docker containers, came the explosion of interest in developing and deploying containerized applications. As the number of containerized applications grew to span hundreds of containers deployed across multiple servers, operating them became more complex. How do you coordinate, scale, manage, and schedule hundreds of containers? This is where Kubernetes can help. Kubernetes is an open source orchestration system that allows you to run your Docker containers and workloads. It helps you manage the operating complexities when moving to scale multiple containers deployed across multiple servers. The Kubernetes engine automatically orchestrates the container lifecycle, distributing the application containers across the hosting infrastructure. Kubernetes can quickly scale resources up or down, depending on the demand. It continually provisions, schedules, deletes, and monitors the health of the containers.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="743 516 1818 883"> <div style="display: flex; justify-content: space-around;"> <div style="text-align: center;"> <p>VMs</p> </div> <div style="text-align: center;"> <p>Containers</p> </div> </div> </div> <div style="display: flex; justify-content: space-between; margin-top: 20px;"> <div data-bbox="835 919 1197 1010"> <p>Virtual Machines</p> <ul style="list-style-type: none"> • Each virtual machine (VM) includes the app, the necessary binaries and libraries and an <u>entire guest operating system</u> </div> <div data-bbox="1276 919 1684 1094"> <p>Containers</p> <ul style="list-style-type: none"> • Containers include the app and all of its dependencies, but <u>share the kernel</u> with other containers. • Run as an isolated process in userspace on the host operating system. • <u>Not</u> tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. </div> </div> <p style="margin-top: 20px;">https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container Cloud Services</p> <p>The first part of this article explained some important Docker concepts. However, in a production environment it is not enough to simply run an application in a Docker container.</p> <p>To setup and operate a production environment requires hardware to run the containers. Software such as Docker, along with repositories and cluster managers, must be installed, upgraded and patched. If several Docker containers communicate across hosts, a network must be created. Clustered containers should be restarted if they fail. In addition, a set of containers linked to each other should be deployable as easily as a single logical application instance. An example of this could be a load balancer, a few web servers, some Oracle WebLogic Server instances with an admin server, a managed server, and a database. To manage containerized applications at scale, requires a container orchestration system like Kubernetes or Docker Swarm. Deploying, managing, and operating orchestration systems like Kubernetes can be challenging and time-consuming.</p> <p>To make it easier and more efficient for developers to create containerized applications, cloud providers offer Container Cloud Services or Containers as a Service (CaaS). Container Cloud Services help developers and operations teams streamline and manage the lifecycle of containers in an automated fashion. These orchestration services, typically built using Kubernetes, make it easier for DevOps teams to manage and operate containerized applications at scale. Oracle Cloud Infrastructure Kubernetes Engine and Azure Kubernetes Service are two examples of popular container orchestration managed cloud services.</p> <p>Oracle Cloud Infrastructure Kubernetes Engine is a fully managed, scalable, and highly available service that you can use to deploy your containerized applications in the cloud. Use Kubernetes Engine (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud native applications.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p>https://www.techtarget.com/searchdatacenter/definition/kernel</p>

Claim 1	Accused Instrumentalities
<p>[1a] storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers;</p>	<p>The method practiced by Oracle and/or its customer through the Accused Instrumentalities includes a step of storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers.</p> <p>For example, OCI and/or OKE stores application containers, sometimes called Docker containers, container images, Kubernetes containers, or Kubernetes pods, in persistent storage available to each node running the application. The terms “node” and “host” are both used to refer to the claimed server. The container might be in a format defined by the Open Container Initiative. This storage may be physically attached to the server or connected through any supported interconnect, including over a network. In addition to the application software, each container includes associated system files, including a Linux user space required to execute the application, for example libc/glibc and other shared libraries, configuration files, etc. necessary for the application. For example, the container includes a base OS image provided by Oracle or by a third party. The container is compatible with the host kernel, for example because the container libraries are linked against the Linux kernel, and the supported host operating systems also use the Linux kernel, which has a stable binary interface.</p> <p>The containers are secure containers as claimed. For example, the data within an individual container is insulated from the effects of other containers except to the extent the container is specifically configured to allow other containers to modify its data, for example using a shared volume.</p> <p><i>See, e.g.:</i></p> <h2>Overview of File Storage</h2> <p>Oracle Cloud Infrastructure File Storage service provides a durable, scalable, secure, enterprise-grade network file system. You can connect to a File Storage service file system from any bare metal, virtual machine, or container instance in your Virtual Cloud Network (VCN). You can also access a file system from outside the VCN using VCN peering, Oracle Cloud Infrastructure FastConnect, and Internet Protocol security (IPSec) virtual private network (VPN).</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="684 207 1016 253">What is Docker?</h2> <hr data-bbox="684 293 749 302"/> <p data-bbox="684 363 1944 521">A Docker container is a packaging format that packages all the code and dependencies of an application in a standard format that allows it to run quickly and reliably across computing environments. A Docker container is a popular lightweight, standalone, executable container that includes everything needed to run an application, including libraries, system tools, code, and runtime. Docker is also a software platform that allows developers to build, test, and deploy containerized applications quickly.</p> <p data-bbox="684 553 1959 646">Containers as a Service (CaaS) or Container Services are managed cloud services that manage the lifecycle of containers. Container services help orchestrate (start, stop, scale) the runtime of containers. Using container services, you can simplify, automate, and accelerate your application development and deployment lifecycle.</p> <p data-bbox="684 678 1902 771">Docker and Container Services have seen rapid adoption and have been a tremendous success over the last several years. From an almost unknown and rather technical open source technology in 2013, Docker has evolved into a standardized runtime environment now officially supported for many Oracle enterprise products.</p> <p data-bbox="674 795 1707 824">https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>Docker versus Kubernetes</p> <p>Linux containers have existed since 2008, but they were not well known until the emergence of Docker containers in 2013. With the onset of Docker containers, came the explosion of interest in developing and deploying containerized applications. As the number of containerized applications grew to span hundreds of containers deployed across multiple servers, operating them became more complex. How do you coordinate, scale, manage, and schedule hundreds of containers? This is where Kubernetes can help. Kubernetes is an open source orchestration system that allows you to run your Docker containers and workloads. It helps you manage the operating complexities when moving to scale multiple containers deployed across multiple servers. The Kubernetes engine automatically orchestrates the container lifecycle, distributing the application containers across the hosting infrastructure. Kubernetes can quickly scale resources up or down, depending on the demand. It continually provisions, schedules, deletes, and monitors the health of the containers.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="743 516 1818 883"> </div> <div data-bbox="835 919 1197 1010"> <p>Virtual Machines</p> <ul style="list-style-type: none"> Each virtual machine (VM) includes the app, the necessary binaries and libraries and an entire guest operating system </div> <div data-bbox="1276 919 1684 1094"> <p>Containers</p> <ul style="list-style-type: none"> Containers include the app and all of its dependencies, but share the kernel with other containers. Run as an isolated process in userspace on the host operating system. Not tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. </div> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker</p> <p>Images and Containers</p> <p>An image is a read-only template with instructions for creating a Docker container and an image is based on another image.</p> <p>A container is a standard unit of software that packages up code and all its dependencies. Hence, the application runs quickly and reliably from one environment to another.</p> <p>A Docker Container Image is a lightweight, standalone, executable package of software that includes everything needed to run an application such as code, runtime, system tools, system libraries, and settings.</p> <p>Container images become containers at runtime and for Docker containers, the images become containers when they run on the engine. Containers are available for both Linux and Windows-based applications. The containerized software always runs the same code, regardless of the infrastructure. The container isolates software from its environment and ensures that it works uniformly despite differences for instance between Development, Staging, and Production.</p> <p>Kubernetes (K8)</p> <p>Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications. It groups the containers that makes an application into logical units for easy management and discovery.</p> <p>https://docs.oracle.com/en/industries/financial-services/microservices-common/14.6.1.0.0/contg/technologies.html</p> <p>Container images</p> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="684 201 1451 224">6. Do Docker containers package up the entire OS and make it easier to deploy?</p> <p data-bbox="684 261 1709 345">Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p data-bbox="674 362 1841 391">https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</p> <h2 data-bbox="684 435 1312 500">About storage drivers</h2> <p data-bbox="684 548 1911 670">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="684 740 1604 797">Storage drivers versus Docker volumes</h2> <p data-bbox="684 834 1953 1101">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="684 1149 1944 1271">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="674 1304 1266 1333">https://docs.docker.com/storage/storagedriver/</p>

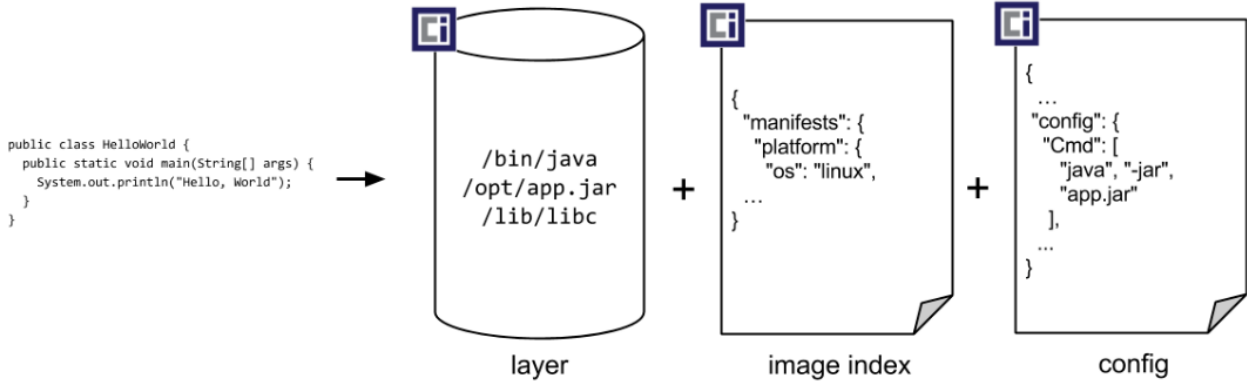
Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 207 1121 261">Images and layers</h2> <p data-bbox="695 302 1860 375">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="695 418 1493 764"># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="695 824 1940 1130">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="674 1154 1266 1187">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p> <div data-bbox="951 706 1707 1299" data-label="Diagram"> <p>The diagram illustrates the Docker layer architecture. It shows a stack of four read-only (R/O) image layers, each represented by a blue box with a unique ID and its size. From bottom to top, the layers are: <code>d3a1f33e8a5a</code> (188.1 MB), <code>c22013c84729</code> (194.5 KB), <code>d74508fb6632</code> (1.895 KB), and <code>91e54dfb1179</code> (0 B). These layers are collectively labeled as 'Image Layers (R/O)' with a padlock icon indicating they are read-only. Above this stack is a dashed box labeled 'Thin R/W layer', which is identified as the 'Container layer' by an arrow. Bidirectional arrows connect the container layer to the top of the image layers, indicating interaction. The entire stack is labeled 'Container (based on ubuntu:15.04 image)' at the bottom.</p> </div> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="690 220 959 282">Volumes</h2> <p data-bbox="690 337 1944 467">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="672 490 1346 521">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="697 573 1264 621">Container environment</h2> <p data-bbox="697 659 1514 724">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="735 763 1488 922" style="list-style-type: none"> • A filesystem, which is a combination of an image and one or more volumes. • Information about the Container itself. • Information about other objects in the cluster. <p data-bbox="672 958 1566 989">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="699 215 915 277">Images</h2> <p data-bbox="699 310 1562 461">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="699 500 1568 570">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="674 597 1367 630">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="695 675 957 737">Volumes</h2> <p data-bbox="695 773 1570 1214">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="674 1242 1346 1274">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="711 220 1337 277">Open Container Initiative</h2> <hr data-bbox="711 289 1948 295"/> <h3 data-bbox="711 342 1222 386">Image Format Specification</h3> <hr data-bbox="711 397 1948 404"/> <p data-bbox="711 435 1948 508">This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p data-bbox="711 544 1948 617">The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p data-bbox="669 646 1520 719">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<p>Overview</p> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 207 1339 267">OCI Image Configuration</h2> <p data-bbox="688 321 1957 483">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="688 519 1701 555">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="672 584 1545 657">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="701 215 789 251">Layer</p> <ul data-bbox="730 289 1955 634" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="701 686 898 722">Image JSON</p> <ul data-bbox="730 760 1955 1105" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="674 1133 1541 1203"> https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md </p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> • rootfs <i>object</i>, REQUIRED <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none"> ◦ type <i>string</i>, REQUIRED <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter a unknown value while verifying or unpacking an image.</p> <ul style="list-style-type: none"> ◦ diff_ids <i>array of strings</i>, REQUIRED <p>An array of layer content hashes (<code>DiffIDs</code>), in order from first to last.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
<p>[1b] wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems,</p>	<p>In the method practiced by Oracle and/or its customer through the Accused Instrumentalities, the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems.</p> <p>The associated system files in the container are compatible with the host kernel, for example because they are linked against the Linux kernel and the supported host operating systems also use the Linux kernel, which has a stable binary interface.</p> <p><i>See discussion in element [1a] above.</i></p> <p><i>See also, e.g.:</i></p> <p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
<p>[1c] the containers of application software excluding a kernel,</p>	<p>In the method practiced by Oracle and/or its customer through the Accused Instrumentalities, the containers of application software exclude a kernel.</p> <p><i>See, e.g.:</i></p> <p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>6. Do Docker containers package up the entire OS and make it easier to deploy?</p> <p>Docker containers do not package up the OS. They package up the applications with everything that the application needs to run. The engine is installed on top of the OS running on a host. Containers share the OS kernel allowing a single host to run multiple containers.</p> <p>https://www.docker.com/blog/the-10-most-common-questions-it-admins-ask-about-docker/</p>

Claim 1	Accused Instrumentalities
<p>[1d] wherein some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by Oracle and/or its customer through the Accused Instrumentalities, some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server.</p> <p>For example, each container will utilize its own associated system files, including libraries such as libc/glibc and configuration files, not the corresponding associated local system files (<i>e.g.</i>, libraries and configuration files of the host OS). As described above and below, in the Accused Instrumentalities the associated system files provide at least some of the same functionalities as the associated local system files. The host/node's associated local system files remain resident on the host/node, for example for use by system processes or applications outside the container environment.</p> <p><i>See, e.g.:</i></p> <p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system's kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker versus Kubernetes</p> <p>Linux containers have existed since 2008, but they were not well known until the emergence of Docker containers in 2013. With the onset of Docker containers, came the explosion of interest in developing and deploying containerized applications. As the number of containerized applications grew to span hundreds of containers deployed across multiple servers, operating them became more complex. How do you coordinate, scale, manage, and schedule hundreds of containers? This is where Kubernetes can help. Kubernetes is an open source orchestration system that allows you to run your Docker containers and workloads. It helps you manage the operating complexities when moving to scale multiple containers deployed across multiple servers. The Kubernetes engine automatically orchestrates the container lifecycle, distributing the application containers across the hosting infrastructure. Kubernetes can quickly scale resources up or down, depending on the demand. It continually provisions, schedules, deletes, and monitors the health of the containers.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="743 516 1818 883"> </div> <div data-bbox="835 919 1197 1010"> <p>Virtual Machines</p> <ul style="list-style-type: none"> Each virtual machine (VM) includes the app, the necessary binaries and libraries and an entire guest operating system </div> <div data-bbox="1276 919 1684 1094"> <p>Containers</p> <ul style="list-style-type: none"> Containers include the app and all of its dependencies, but share the kernel with other containers. Run as an isolated process in userspace on the host operating system. Not tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. </div> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container Cloud Services</p> <p>The first part of this article explained some important Docker concepts. However, in a production environment it is not enough to simply run an application in a Docker container.</p> <p>To setup and operate a production environment requires hardware to run the containers. Software such as Docker, along with repositories and cluster managers, must be installed, upgraded and patched. If several Docker containers communicate across hosts, a network must be created. Clustered containers should be restarted if they fail. In addition, a set of containers linked to each other should be deployable as easily as a single logical application instance. An example of this could be a load balancer, a few web servers, some Oracle WebLogic Server instances with an admin server, a managed server, and a database. To manage containerized applications at scale, requires a container orchestration system like Kubernetes or Docker Swarm. Deploying, managing, and operating orchestration systems like Kubernetes can be challenging and time-consuming.</p> <p>To make it easier and more efficient for developers to create containerized applications, cloud providers offer Container Cloud Services or Containers as a Service (CaaS). Container Cloud Services help developers and operations teams streamline and manage the lifecycle of containers in an automated fashion. These orchestration services, typically built using Kubernetes, make it easier for DevOps teams to manage and operate containerized applications at scale. Oracle Cloud Infrastructure Kubernetes Engine and Azure Kubernetes Service are two examples of popular container orchestration managed cloud services.</p> <p>Oracle Cloud Infrastructure Kubernetes Engine is a fully managed, scalable, and highly available service that you can use to deploy your containerized applications in the cloud. Use Kubernetes Engine (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud native applications.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
<p>[1e] wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by Oracle and/or its customer through the Accused Instrumentalities, said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server.</p> <p>For example, in some cases the host OS and container will use one or more identical system files, for example when both the host and the container incorporate the same Linux distribution version, or when both host and container use the same version of libc. In the case where the associated system files are identical to the associated local system files, they are copies thereof. In other cases modified copies are used instead, for example when different versions of the same library, or configuration files with different parameters, are used by the host and container.</p> <p><i>See, e.g.:</i></p> <p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker versus Kubernetes</p> <p>Linux containers have existed since 2008, but they were not well known until the emergence of Docker containers in 2013. With the onset of Docker containers, came the explosion of interest in developing and deploying containerized applications. As the number of containerized applications grew to span hundreds of containers deployed across multiple servers, operating them became more complex. How do you coordinate, scale, manage, and schedule hundreds of containers? This is where Kubernetes can help. Kubernetes is an open source orchestration system that allows you to run your Docker containers and workloads. It helps you manage the operating complexities when moving to scale multiple containers deployed across multiple servers. The Kubernetes engine automatically orchestrates the container lifecycle, distributing the application containers across the hosting infrastructure. Kubernetes can quickly scale resources up or down, depending on the demand. It continually provisions, schedules, deletes, and monitors the health of the containers.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="745 516 1820 881"> </div> <div data-bbox="835 917 982 938"> <p>Virtual Machines</p> </div> <div data-bbox="835 951 1197 1006"> <ul style="list-style-type: none"> Each virtual machine (VM) includes the app, the necessary binaries and libraries and an entire guest operating system </div> <div data-bbox="1276 917 1375 938"> <p>Containers</p> </div> <div data-bbox="1276 951 1682 1091"> <ul style="list-style-type: none"> Containers include the app and all of its dependencies, but share the kernel with other containers. Run as an isolated process in userspace on the host operating system. Not tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. </div> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container Cloud Services</p> <p>The first part of this article explained some important Docker concepts. However, in a production environment it is not enough to simply run an application in a Docker container.</p> <p>To setup and operate a production environment requires hardware to run the containers. Software such as Docker, along with repositories and cluster managers, must be installed, upgraded and patched. If several Docker containers communicate across hosts, a network must be created. Clustered containers should be restarted if they fail. In addition, a set of containers linked to each other should be deployable as easily as a single logical application instance. An example of this could be a load balancer, a few web servers, some Oracle WebLogic Server instances with an admin server, a managed server, and a database. To manage containerized applications at scale, requires a container orchestration system like Kubernetes or Docker Swarm. Deploying, managing, and operating orchestration systems like Kubernetes can be challenging and time-consuming.</p> <p>To make it easier and more efficient for developers to create containerized applications, cloud providers offer Container Cloud Services or Containers as a Service (CaaS). Container Cloud Services help developers and operations teams streamline and manage the lifecycle of containers in an automated fashion. These orchestration services, typically built using Kubernetes, make it easier for DevOps teams to manage and operate containerized applications at scale. Oracle Cloud Infrastructure Kubernetes Engine and Azure Kubernetes Service are two examples of popular container orchestration managed cloud services.</p> <p>Oracle Cloud Infrastructure Kubernetes Engine is a fully managed, scalable, and highly available service that you can use to deploy your containerized applications in the cloud. Use Kubernetes Engine (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud native applications.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>COPY and ADD : These commands copy files and directories from your local filesystem into the Docker image. They are often used to include your application code, configuration files, and dependencies.</p> <p>https://medium.com/@swalperen3008/what-is-dockerize-and-dockerize-your-project-a-step-by-step-guide-899c48a34df6</p>

Claim 1	Accused Instrumentalities
	<p>Container images</p> <p>A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p>https://kubernetes.io/docs/concepts/containers/</p>

Claim 1	Accused Instrumentalities
<p>[1f] and wherein the application software cannot be shared between the plurality of secure containers of application software,</p>	<p>In the method practiced by Oracle and/or its customer through the Accused Instrumentalities, the application software cannot be shared between the plurality of secure containers of application software.</p> <p>For example, each container has an isolated runtime environment that cannot be accessed by other containers, for example including a per-container writeable layer or other ephemeral per-container storage. For another example, when the plurality of secure containers each corresponds to a different container image, each container cannot access another container's image and therefore application software.</p> <p><i>See, e.g.:</i></p> <p>Cgroups and Namespaces History</p> <p>The underlying Linux kernel features that Docker uses are cgroups and namespaces. In 2008 cgroups were introduced to the Linux kernel based on work previously done by Google developers¹. Cgroups limit and account for the resource usage of a set of operating system processes.</p> <p>The Linux kernel uses namespace to isolate the system resources of processes from each other. The first namespace, i.e. the mount namespace, was introduced as early as 2002.²</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Container:</p> <p>Unlike a VM which provides hardware virtualization, a container provides lightweight, operating-system-level virtualization by abstracting the “user space.” Containers share the host system’s kernel with other containers. A container, which runs on the host operating system, is a standard software unit that packages code and all its dependencies, so applications can run quickly and reliably from one environment to another. Containers are nonpersistent and are spun up from images.</p> <p>Docker engine:</p> <p>The open source host software building and running the containers. Docker Engines act as the client-server application supporting containers on various Windows servers and Linux operating systems, including Oracle Linux, CentOS, Debian, Fedora, RHEL, SUSE, and Ubuntu.</p> <p>Docker images:</p> <p>Collection of software to be run as a container that contains a set of instructions for creating a container that can run on the Docker platform. Images are immutable, and changes to an image require to build a new image.</p> <p>Docker Registry:</p> <p>Place to store and download images. The registry is a stateless and scalable server-side application that stores and distributes Docker images.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p> <p>Docker versus Kubernetes</p> <p>Linux containers have existed since 2008, but they were not well known until the emergence of Docker containers in 2013. With the onset of Docker containers, came the explosion of interest in developing and deploying containerized applications. As the number of containerized applications grew to span hundreds of containers deployed across multiple servers, operating them became more complex. How do you coordinate, scale, manage, and schedule hundreds of containers? This is where Kubernetes can help. Kubernetes is an open source orchestration system that allows you to run your Docker containers and workloads. It helps you manage the operating complexities when moving to scale multiple containers deployed across multiple servers. The Kubernetes engine automatically orchestrates the container lifecycle, distributing the application containers across the hosting infrastructure. Kubernetes can quickly scale resources up or down, depending on the demand. It continually provisions, schedules, deletes, and monitors the health of the containers.</p> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<p>Docker Basics</p> <p>The core concepts of Docker are images and containers. A Docker image contains everything that is needed to run your software: the code, a runtime (for example, Java Virtual Machine (JVM), drivers, tools, scripts, libraries, deployments, and more.</p> <p>A Docker container is a running instance of a Docker image. However, unlike in traditional virtualization with a type 1 or type 2 hypervisor, a Docker container runs on the kernel of the host operating system. Within a Docker image there is no separate operating system, as illustrated in Figure 1.</p> <div data-bbox="745 516 1820 881"> </div> <div data-bbox="835 917 982 938"> <p>Virtual Machines</p> </div> <div data-bbox="835 954 1199 1008"> <ul style="list-style-type: none"> Each virtual machine (VM) includes the app, the necessary binaries and libraries and an entire guest operating system </div> <div data-bbox="1281 917 1375 938"> <p>Containers</p> </div> <div data-bbox="1281 954 1682 1092"> <ul style="list-style-type: none"> Containers include the app and all of its dependencies, but share the kernel with other containers. Run as an isolated process in userspace on the host operating system. Not tied to any specific infrastructure - containers run on any computer, on any infrastructure and in any cloud. </div> <p>https://www.oracle.com/in/cloud/cloud-native/container-registry/what-is-docker/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="682 212 1312 277">About storage drivers</h2> <p data-bbox="682 326 1913 448">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 518 1604 574">Storage drivers versus Docker volumes</h2> <p data-bbox="682 613 1953 878">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="682 930 1944 1052">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="672 1084 1266 1117">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="695 207 1121 261">Images and layers</h2> <p data-bbox="695 302 1860 375">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="695 418 1493 764"># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="695 824 1940 1130">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="674 1154 1266 1187">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p> <div data-bbox="945 706 1711 1299"> <p>The diagram illustrates the layer structure of a Docker container. At the bottom is a box labeled 'Container (based on ubuntu:15.04 image)'. Inside this box are four stacked blue rectangles representing image layers, each with a hash and a size: '91e54dfb1179' (0 B), 'd74508fb6632' (1.895 KB), 'c22013c84729' (194.5 KB), and 'd3a1f33e8a5a' (188.1 MB). To the right of these layers is a padlock icon and the text 'Image Layers (R/O)'. Above the container box is a dashed rectangle labeled 'Thin R/W layer'. An arrow points from the text 'Container layer' to this dashed box. Vertical double-headed arrows connect the 'Thin R/W layer' to each of the four image layers below it.</p> </div> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
<p>[1g] and wherein each of the containers has a unique root file system that is different from an operating system's root file system.</p>	<p>In the method practiced by Oracle and/or its customer through the Accused Instrumentalities, each of the containers has a unique root file system that is different from an operating system's root file system.</p> <p>For example, the container's root file system comprises the image layer(s), an ephemeral writeable layer (e.g., in Docker terminology the container layer), and optionally one or more volumes. This root file system is distinct and isolated from the host operating system's root file system.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="709 215 1770 269">Setting Up Storage for Kubernetes Clusters</h2> <p data-bbox="709 313 1955 456"><i>Find out how to define and apply persistent volume claims to clusters you've created using Kubernetes Engine (OKE). With Oracle Cloud Infrastructure as the underlying IaaS provider, you can provision persistent volume claims by attaching volumes from the Block Volume service or by mounting file systems from the File Storage service.</i></p> <p data-bbox="709 500 1917 605">Container storage via a container's root file system is ephemeral, and can disappear upon container deletion and creation. To provide a durable location to prevent data from being lost, you can create and use persistent volumes to store data outside of containers.</p> <p data-bbox="709 649 1866 716">A persistent volume offers persistent storage that enables your data to remain intact, regardless of whether the containers to which the storage is connected are terminated.</p> <p data-bbox="709 760 1934 826">A persistent volume claim (PVC) is a request for storage, which is met by binding the PVC to a persistent volume (PV). A PVC provides an abstraction layer to the underlying storage.</p> <p data-bbox="709 870 1614 899">With Oracle Cloud Infrastructure, you can provision persistent volume claims:</p> <ul data-bbox="737 943 1925 1300" style="list-style-type: none"> • By attaching volumes from the Oracle Cloud Infrastructure Block Volume service. The volumes are connected to clusters created by Kubernetes Engine using CSI (Container Storage Interface) or FlexVolume volume plugins deployed on the clusters. Oracle recommends the CSI volume plugin since the upstream Kubernetes project deprecates the FlexVolume volume plugin in Kubernetes version 1.23. See Provisioning PVCs on the Block Volume Service. • By mounting file systems in the Oracle Cloud Infrastructure File Storage service. The File Storage service file systems are mounted inside containers running on clusters created by Kubernetes Engine using a CSI (Container Storage Interface) volume plugin deployed on the clusters. See Provisioning PVCs on the File Storage Service. <p data-bbox="674 1312 1633 1380">https://docs.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingpersistentvolumeclaim.htm</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="695 212 1178 251">Configuring Docker Storage</p> <p data-bbox="695 282 1961 423">The Docker Engine is configured to use <code>overlay2</code> as the default storage driver to manage Docker containers. This provides a performance and scalability improvement on earlier releases that used the device mapper as the default storage driver, but the technology is new and should be tested properly before use in production environments. For more information on <code>overlay2</code>, see:</p> <p data-bbox="695 461 1570 488">https://docs.docker.com/engine/userguide/storagedriver/overlayfs-driver/</p> <p data-bbox="695 521 1898 586">Overlay file systems can corrupt when used in conjunction with any file system that does not have dtype support enabled.</p> <p data-bbox="674 613 1612 678">https://docs.oracle.com/en/operating-systems/oracle-linux/docker/docker-InstallingOracleContainerRuntimeforDocker.html</p> <p data-bbox="690 711 1944 906">In Kubernetes, each container can read and write to its own file system. But when a container is restarted, all data is lost. Therefore, containers that need to maintain state would store data in a persistent storage such as Network File System (NFS). What's already stored in NFS isn't deleted when a pod, which might contain one or more containers, is destroyed. Also, an NFS can be accessed from multiple pods at the same time, so an NFS can be used to share data between pods. This behavior is really useful when containers or applications need to read configuration data from a single shared file system or when multiple containers need to read from and write data to a single shared file system.</p> <p data-bbox="690 938 1974 1170">Oracle Cloud Infrastructure File Storage provides a durable, scalable, and distributed enterprise-grade network file system that supports NFS version 3 along with Network Lock Manager (NLM) for a locking mechanism. You can connect to File Storage from any bare metal, virtual machine, or container instance in your virtual cloud network (VCN). You can also access a file system from outside the VCN by using Oracle Cloud Infrastructure FastConnect or an Internet Protocol Security (IPSec) virtual private network (VPN). File Storage is a fully managed service so you don't have to worry about hardware installations and maintenance, capacity planning, software upgrades, security patches, and so on. You can start with a file system that contains only a few kilobytes of data and grows to handle 8 exabytes of data.</p> <p data-bbox="674 1187 1940 1252">https://blogs.oracle.com/cloud-infrastructure/post/using-file-storage-service-with-container-engine-for-kubernetes</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="682 212 1312 277">About storage drivers</h2> <p data-bbox="682 326 1913 448">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="682 518 1604 574">Storage drivers versus Docker volumes</h2> <p data-bbox="682 612 1953 878">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="682 927 1944 1052">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1381 976 1575 1003" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="674 1084 1266 1117"><a data-bbox="674 1084 1266 1117" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

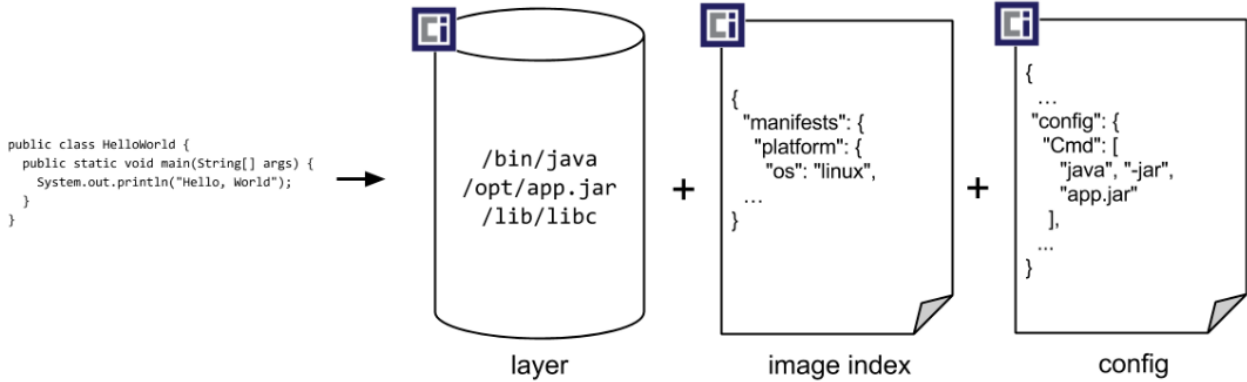
Claim 1	Accused Instrumentalities
	<h2 data-bbox="699 207 1119 264">Images and layers</h2> <p data-bbox="699 302 1860 378">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="699 418 1946 784"># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p data-bbox="699 824 1938 1133">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The FROM statement starts out by creating a layer from the ubuntu:22.04 image. The LABEL command only modifies the image's metadata, and doesn't produce a new layer. The COPY command adds some files from your Docker client's current directory. The first RUN command builds your application using the make command, and writes the result to a new layer. The second RUN command removes a cache directory, and writes the result to a new layer. Finally, the CMD instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="674 1154 1266 1187">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p>Each layer is only a set of differences from the layer before it. Note that both <i>adding</i>, and <i>removing</i> files will result in a new layer. In the example above, the <code>\$HOME/.cache</code> directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the Best practices for writing Dockerfiles and use multi-stage builds sections to learn how to optimize your Dockerfiles for efficient images.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p> <div data-bbox="953 704 1709 1305"> <p>The diagram illustrates the Docker layer architecture. It shows a stack of four read-only (R/O) image layers, each represented by a blue box with a unique ID and its size. From bottom to top, the layers are: <code>d3a1f33e8a5a</code> (188.1 MB), <code>c22013c84729</code> (194.5 KB), <code>d74508fb6632</code> (1.895 KB), and <code>91e54dfb1179</code> (0 B). These layers are collectively labeled as 'Image Layers (R/O)' with a padlock icon indicating they are read-only. Above this stack is a dashed box labeled 'Thin R/W layer', which is identified as the 'Container layer' by an arrow. Bidirectional arrows connect the container layer to the top of the image layers. The entire stack is labeled 'ubuntu:15.04' and 'Container (based on ubuntu:15.04 image)'.</p> </div> <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="690 220 959 282">Volumes</h2> <p data-bbox="690 337 1944 467">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="672 490 1346 521">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="697 573 1264 621">Container environment</h2> <p data-bbox="697 660 1514 724">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="735 763 1488 922" style="list-style-type: none"> • A filesystem, which is a combination of an image and one or more volumes. • Information about the Container itself. • Information about other objects in the cluster. <p data-bbox="672 959 1566 990">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="699 215 915 277">Images</h2> <p data-bbox="699 310 1562 461">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="699 500 1568 570">You typically create a container image of your application and push it to a registry before referring to it in a Pod.</p> <p data-bbox="674 597 1367 630">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="695 675 957 737">Volumes</h2> <p data-bbox="695 773 1570 1214">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a Pod and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes volume abstraction solves both of these problems. Familiarity with Pods is suggested.</p> <p data-bbox="674 1242 1346 1274">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="711 219 1337 277">Open Container Initiative</h2> <hr data-bbox="711 289 1948 295"/> <h3 data-bbox="711 342 1222 386">Image Format Specification</h3> <hr data-bbox="711 397 1948 404"/> <p data-bbox="711 435 1948 508">This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p data-bbox="711 544 1948 617">The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p data-bbox="669 646 1520 719">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<p>Overview</p> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="688 207 1339 267">OCI Image Configuration</h2> <p data-bbox="688 321 1957 483">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="688 519 1701 555">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="672 584 1545 657">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p data-bbox="701 215 789 251">Layer</p> <ul data-bbox="730 289 1955 634" style="list-style-type: none">• Image filesystems are composed of <i>layers</i>.• Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer.• Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer.• Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="701 686 898 722">Image JSON</p> <ul data-bbox="730 760 1955 1105" style="list-style-type: none">• Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes.• The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers.• This JSON is considered to be immutable, because changing it would change the computed ImageID.• Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="674 1133 1541 1203">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> • rootfs <i>object</i>, REQUIRED <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none"> ◦ type <i>string</i>, REQUIRED <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter a unknown value while verifying or unpacking an image.</p> <ul style="list-style-type: none"> ◦ diff_ids <i>array of strings</i>, REQUIRED <p>An array of layer content hashes (<code>DiffIDs</code>), in order from first to last.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>